Proceedings of the 2016 Industrial and Systems Engineering Research Conference
H. Yang, Z. Kong, and MD Sarder, eds.

arXiv:1604.01075v1 [stat.ML] 4 Apr 2016

A Dynamic Bayesian Network Model for Inventory Level Estimation in Retail Marketing

Luis I. Reyes-Castro **Andres G. Abad**
Escuela Superior Politécnica del Litoral (ESPOL)
Guayaquil - Ecuador

Abstract

Many retailers today employ inventory management systems based on Re-Order Point Policies, most of which rely on the assumption that all decreases in product inventory levels result from product sales. Unfortunately, it usually happens that small but random quantities of the product get lost, stolen or broken without record as time passes, *e.g.*, as a consequence of shoplifting. This is usual for retailers handling large varieties of inexpensive products, *e.g.*, grocery stores. In turn, over time these discrepancies lead to *stock freezing* problems (see Ref. [1]), *i.e.*, situations where the system believes the stock is above the re-order point but the actual stock is at zero, and so no replenishments or sales occur. Motivated by these issues, we model the interaction between sales, losses, replenishments and inventory levels as a Dynamic Bayesian Network (DBN), where the inventory levels are unobserved (*i.e.*, hidden) variables we wish to estimate. We present an Expectation-Maximization (EM) algorithm to estimate the parameters of the sale and loss distributions, which relies on solving a one-dimensional dynamic program for the E-step and on solving two separate one-dimensional nonlinear programs for the M-step.

Keywords

Inventory Management, Inventory Record Inaccuracies, Inventory Shrinkage, Dynamic Bayesian Networks.

1. Introduction

We consider the sale and replenishment of a product in a store over a time horizon of $T \geq 2$ time periods. We let I_0 denote the initial inventory level, and for each period $t \in \{1, \ldots, T\}$ we let I_t denote the inventory level at the end of that period. Furthermore, for each period t we let S_t denote the random but observed number of units of the product sold during that period. Moreover, we assume that replenishments happen after all the sales of the period have been completed, *e.g.*, after the store closes for the day, and for each period t we let R_t denote the non-random and observed number of units of the product replenished at the end of that period. Finally, and most importantly, we assume that on each time period some number of units of the product may be lost, broken or stolen without knowledge of the store's manager, *i.e.*, without record. In particular, for each period t we let L_t denote the random and unobserved number of units of the product lost, broken or stolen during that period. For modeling reasons, we further assume that in each period all losses occur after all sales have been completed but before any replenishments arrive, although in reality sale, loss and replenishment epochs may intertwine.

Since the product losses are unobserved, so are the inventory levels, which motivates the main problem of this paper: Estimating the (unknown) sale and loss distribution parameters along with the (unobserved) inventory levels. This problem is important because having a good model of the inventory level history of a product is essencial to knowing when to re-order it so as to keep it available to the customers. Unfortunately, most studies so far have focused on qualitatively describing the problem and on proposing heuristic replenishment and inspection policies; the reader is refered to Refs. [1–6] for a sample of previous work. For instance, Ref. [5], which in our opinion is the study most closely related to our paper, describes a method for estimating the aforementioned parameters by collecting statistics of past inventory inspection data and pooling the statistics associated with similar products. In contrast, we propose an algorithm capable of estimating the parameters even in the absence of past inspection data.

2. Assumptions, Problem Statement, and Solution Method

As usual, we assume that for each time period the physical inventory level at the end of the period is equal to the physical inventory level at the end of the previous period, minus the product sales and losses during the period, plus the replenishments at the end of the period. *I.e.:*

$$\forall\, t \in \{\, 1,\, \ldots,\, T \}\, :\; I_t \;=\; I_{t-1} - S_t - L_t + R_t \tag{1}$$

Furthermore, for each period t we assume that S_t has a truncated Poisson distribution with parameter $\sigma > 0$ and upper bound I_{t-1}. *I.e.,* if $(X_t)_{1 \leq t \leq T}$ is a sequence of i.i.d. Poisson random variables with parameter σ then for each t we have $S_t = \min \{\, X_t,\, I_{t-1} \,\}$. We choose the Poisson distribution because it is commonly used to model random demands, although it is fairly straighforward to extend our model to one with a different sales distribution, *e.g.,* Bernoulli, Geometric, Binomial, etc. The truncation is justified because in each period the number of units of the product that the store can sell is limited by the product's physical inventory level at the end of the previous period; we do not allow backordering. Moreover, since the value of I_{t-1} is all we need to describe the distribution of S_t, we observe that conditional on I_{t-1} the random variable S_t is independent of all inventory levels up to period $t-2$ and of all sales, losses and replenishments up to period $t-1$. More precisely, for each $t \in \{\, 1,\, \ldots,\, T \}$:

$$\mathbb{P}\,(\, S_t \mid (I_\tau)_{0 \leq \tau \leq t-1},\, (S_\tau)_{1 \leq \tau \leq t-1},\, (L_\tau)_{1 \leq \tau \leq t-1},\, (R_\tau)_{1 \leq \tau \leq t-1} \,) \;=\; \mathbb{P}\,(\, S_t \mid I_{t-1} \,) \tag{2}$$

In addition, for each period t we assume that L_t has a truncated Bernoulli distribution with parameter $\lambda \in [0, 1]$ and upper bound $I_{t-1} - S_t$, *i.e.,* if $(Y_t)_{1 \leq t \leq T}$ is a sequence of i.i.d. Bernoulli random variables with parameter λ then for each t we have $L_t = \min \{\, Y_t,\, I_{t-1} - S_t \,\}$. We choose the Bernoulli distribution because we are interested in modeling small loss rates (*i.e.,* rates of no more than a unit per period), but our model also allows other discrete distributions. The truncation relies on the fact that in each period all losses occur once all sales have been completed but before any replenishments. Also, since the values of I_{t-1} and S_t completely specify the distribution of L_t, we see that conditional on I_{t-1} and S_t the random variable L_t is independent of all inventory levels up to period $t-2$ and of all sales, losses and replenishments up to period $t-1$. More precisely, for each $t \in \{\, 1,\, \ldots,\, T \}$:

$$\mathbb{P}\,(\, L_t \mid (I_\tau)_{0 \leq \tau \leq t-1},\, (S_\tau)_{1 \leq \tau \leq t-1},\, (L_\tau)_{1 \leq \tau \leq t-1},\, (R_\tau)_{1 \leq \tau \leq t-1} \,) \;=\; \mathbb{P}\,(\, L_t \mid I_{t-1},\, S_t \,) \tag{3}$$

Now, we can continue building our model without product loss variables (*i.e.,* L_t's) if we note, from Equation (1), that for each $t \in \{\, 1,\, \ldots,\, T \}$:

$$\mathbb{P}\,(\, I_t \mid I_{t-1},\, S_t,\, R_t \,) \;=\; \mathbb{P}\,(\, L_t = I_{t-1} - S_t + R_t - I_t \mid I_{t-1},\, S_t \,) \tag{4}$$

Furthermore, combining Equations (3) and (4), we see that conditional on I_{t-1}, S_t and R_t the random variable I_t is independent of all inventory levels up to period $t-2$ and of all sales, losses and replenishments up to period $t-1$. Moreover, since our assumptions require that for each period t all sales and losses occur before any replenishments, we see that $S_t + L_t \leq I_{t-1}$, and so replacing I_{t-1} using Equation (1) we obtain the bound $R_t \;\leq\; I_t$. In turn, combining this bound with the same equation and the fact that $L_t \in \{\, 0, 1 \}$ surely, we have:

$$\forall\, t \in \{\, 1,\, \ldots,\, T \}\, :\;\; \max \{\, R_t\,,\, I_{t-1} - S_t + R_t - 1 \,\} \;\leq\; I_t \;\leq\; I_{t-1} - S_t + R_t \tag{5}$$

Finally, the problem we seek to solve is that of finding a Maximum Likelihood Estimate (MLE) of the sale and loss distribution parameters (σ, λ), respectively, and of the unobserved inventory level history $I \triangleq (I_t)_{1 \leq t \leq T}$, given an initial inventory level I_0, a sales history $S \triangleq (S_t)_{1 \leq t \leq T}$, and a replenishment history $R \triangleq (R_t)_{1 \leq t \leq T}$. More precisely, we seek to solve the following Mixed-Integer Nonlinear Program (MINLP):

$$\text{maximize:} \quad \log \mathbb{P}_{\sigma,\lambda}\,(\, I \mid I_0,\, S,\, R \,) \tag{6}$$

$$\text{subject to the constraints:} \quad \sigma > 0;\; \lambda \in [0, 1];\; I \in \mathbb{Z}^T_{\geq 0};\; \text{Inequality (5)} \tag{7}$$

Unfortunately, to the best of our knowledge there are no efficient methods for computing a global maximizer to Problem (6)-(7) with provable guarantees. Therefore, in this work we present an Expectation-Maximization (EM) Algorithm, which is a greedy algorithm, to compute a local maximizer of the aforementioned function (see Ref. [7]). The algorithm itself is quite simple, relying on the following iteration:

1. Select initial guesses for the MLEs of the sale and loss distribution parameters, denoted $(\, \sigma^{[0]},\, \lambda^{[0]} \,)$.
2. For each iteration $k \geq 1$:
   (a) **E-step:** Using the previous iteration's MLEs of the sale and loss parameters, *i.e.,* $(\, \sigma^{[k-1]},\, \lambda^{[k-1]} \,)$, compute the current MLE of the inventory level history, denoted $I^{[k]} \triangleq (I_t)^{[k]}_{1 \leq t \leq T}$.
   (b) **M-step:** Using the current MLE of the inventory level history, *i.e.,* $I^{[k]}$, compute the current MLEs of the sale and loss parameters, denoted $(\sigma^{[k]},\, \lambda^{[k]})$.

(c) If the sale and loss parameters haven't changed, *i.e.*, if $(\sigma^{[k-1]}, \lambda^{[k-1]}) = (\sigma^{[k]}, \lambda^{[k]})$, then terminate.

Execution of the E-step and of the M-step is explained in the following two sections.

3. E-step: Estimation of the Inventory Level History

In this section we seek to compute an inventory level history $I \in \mathbb{Z}^{T}_{\geq 0}$ of maximum log-likelihood, among all those which satisfy Inequality (5), given the sale and loss distribution parameters $\sigma \in \mathbb{R}_{>0}$ and $\lambda \in [0, 1]$ and conditional on the initial inventory level $I_0 \in \mathbb{Z}_{\geq 0}$, the sales history $S \in \mathbb{Z}^{T}_{\geq 0}$, and the replenishment history $R \in \mathbb{Z}^{T}_{\geq 0}$. To start, we note that since I_0, S and R are fixed, so is their joint likelihood, *i.e.*, $\mathbb{P}_{\sigma,\lambda}(I_0, S, R)$. Therefore, instead of seeking an inventory level history of maximum conditional log-likelihood, we can search for one that maximizes the joint log-likelihood of all the variables, *i.e.*, the function $I \mapsto \log \mathbb{P}_{\sigma,\lambda}(I_0, I, S, R)$.

Next, we recognize that the joint likelihood of all the variables can be factored as a Dynamic Bayesian Network (DBN), *i.e.*, a Bayesian Network with a chain-like structure where each type of node represents a time history (see Ref. [8]). Certainly, from the statistical assumptions put forth in Section 2, we can represent the joint likelihood function as a DBN where each variable is a node, and where the nodes' parents are as follows:

- Nodes I_0, R_1, ..., and R_T have no parents.
- For each period $t \in \{1, \ldots, T\}$, the parent of node S_t is node I_{t-1}.
- For each period $t \in \{1, \ldots, T\}$, the parents of node I_t are nodes I_{t-1}, S_t and R_t.

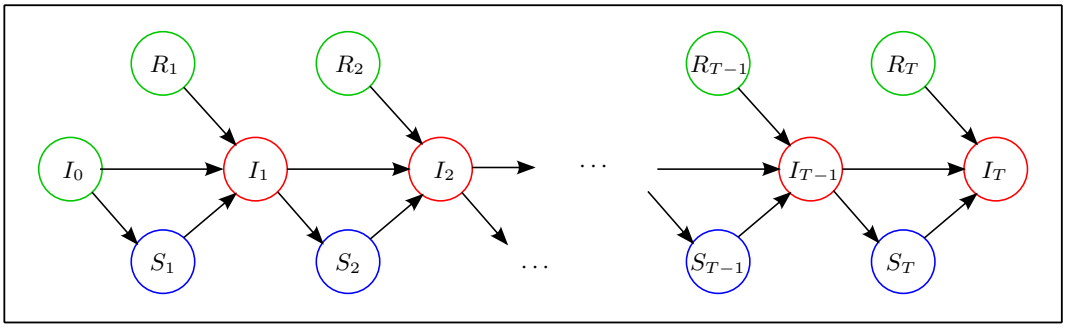

Figure 1: Our Dynamic Bayesian Network (DBN) model. Here, deterministic variable nodes are shown in green, observed random variable nodes are shown in blue, and unobserved random variable nodes are shown in red.

Now, with this representation in mind we can factor the joint log-likelihood function as follows:

$$\begin{aligned} \log \mathbb{P}_{\sigma,\lambda}(I_0, I, S, R) &= \log\left(\prod_{t=1}^{T} \mathbb{P}_{\sigma}(S_t \mid I_{t-1})\, \mathbb{P}_{\lambda}(I_t \mid I_{t-1}, S_t, R_t)\right) \\ &= \sum_{t=1}^{T-1}\left(\log \mathbb{P}_{\sigma}(S_t \mid I_{t-1}) + \log \mathbb{P}_{\lambda}(I_t \mid I_{t-1}, S_t, R_t)\right) \end{aligned} \tag{8}$$

Furthermore, if for each $t \in \{1, \ldots, T\}$ we define the function

$$\phi^{[t]}_{\sigma,\lambda}(I_{t-1}, I_t) \triangleq \log \mathbb{P}_{\sigma}(S_t \mid I_{t-1}) + \log \mathbb{P}_{\lambda}(I_t \mid I_{t-1}, S_t, R_t), \tag{9}$$

then it is clear that our problem is equivalent to that of maximizing the function:

$$\Phi_{\sigma,\lambda}(I) = \sum_{t=1}^{T} \phi^{[t]}_{\sigma,\lambda}(I_{t-1}, I_t) \tag{10}$$

Moreover, the function above can be maximized sequentially by means of a Dynamic Programming (DP) approach. Indeed, notice the following:

$$\max_{I \in \mathbb{Z}^{T}_{\geq 0}} \Phi_{\sigma,\lambda}(I) = \max_{I_1, \ldots, I_{T-1}} \left[\left(\sum_{t=1}^{T-1} \phi^{[t]}_{\sigma,\lambda}(I_{t-1}, I_t)\right) + \underbrace{\max_{I_T \in \mathbb{Z}_{\geq 0}} \phi^{[T]}_{\sigma,\lambda}(I_{T-1}, I_T)}_{\triangleq\ \psi^{[T-1]}_{\sigma,\lambda}(I_{T-1})}\right]$$

$$
\begin{aligned}
&= \max_{I_1,\,\ldots,\,I_{T-2}} \left[\left(\sum_{t=1}^{T-2} \phi_{\sigma,\lambda}^{[t]}(\, I_{t-1},\, I_t\,) \right) + \underbrace{\max_{I_{T-1}\,\in\,\mathbb{Z}_{\geq 0}} \left[\phi_{\sigma,\lambda}^{[T-1]}(\, I_{T-2},\, I_{T-1}\,) + \psi_{\sigma,\lambda}^{[T-1]}(\, I_{T-1}\,) \right]}_{\triangleq\ \psi_{\sigma,\lambda}^{[T-2]}(\, I_{T-2}\,)} \right] \\
&= \cdots \\
&= \max_{I_1,\,\ldots,\,I_{T-k}} \left[\left(\sum_{t=1}^{T-k} \phi_{\sigma,\lambda}^{[t]}(\, I_{t-1},\, I_t\,) \right) + \underbrace{\max_{I_{T-k+1}\,\in\,\mathbb{Z}_{\geq 0}} \left[\phi_{\sigma,\lambda}^{[T-k+1]}(\, I_{T-k},\, I_{T-k+1}\,) + \psi_{\sigma,\lambda}^{[T-k+1]}(\, I_{T-k+1}\,) \right]}_{\triangleq\ \psi_{\sigma,\lambda}^{[T-k]}(\, I_{T-k}\,)} \right]
\end{aligned}
\tag{11}
$$

Hence, in light of the previous arguments, Algorithm 1 computes a sequence of feasible inventory levels $I^* \in \mathbb{Z}_{\geq 0}^T$ which maximizes function $\Phi_{\sigma,\lambda}$ among all feasible inventory level histories.

Algorithm 1: Dynamic Programming Algorithm for the E-step

Data: Sale and loss distribution parameters $\sigma \in \mathbb{R}_{>0}$ and $\lambda \in [0,1]$, initial inventory level $I_0 \in \mathbb{Z}_{\geq 0}$, sales history $S \triangleq (S_t)_{1\leq t\leq T} \in \mathbb{Z}_{\geq 0}^T$, replenishment history $R \triangleq (R_t)_{1\leq t\leq T} \in \mathbb{Z}_{\geq 0}^T$.

Result: Maximum likelihood estimate (MLE) of the inventory level history $I^* \triangleq (I_t^*)_{1\leq t\leq T}$.

`// SETUP: Computes upper bounds on the feasible inventory levels`

let $I_0^{max} = I_0$

foreach $t \in (\,1,\,\ldots,\,T\,)$ **do**

  let $I_t^{max} = I_{t-1}^{max} - S_t + R_t$

`// BACKWARD PASS: Computes the optimal` I_{t+1}`'s as functions of the` I_t`'s`

let $\psi_{\sigma,\lambda}^{[T]}$ be a zero function, *i.e.*, $\psi_{\sigma,\lambda}^{[T]} : \mathbb{Z}_{\geq 0} \mapsto \{0\}$

foreach $t \in (\,T-1,\,\ldots,\,0\,)$ **do**

  foreach $I_t \in \{\,0,\,\ldots,\,I_t^{max}\,\}$ **do**

    let $\psi_{\sigma,\lambda}^{[t]}(\, I_t\,)$ be the maximum with respect to I_{t+1}, and $\omega_{\sigma,\lambda}^{[t]}(\, I_t\,)$ be a maximizing value of I_{t+1}, respectively, of the function:

$$
I_t \quad \mapsto \quad \phi_{\sigma,\lambda}^{[t+1]}(\, I_t,\, I_{t+1}\,) + \psi_{\sigma,\lambda}^{[t+1]}(\, I_{t+1}\,)
$$

    Note: This is a finite maximization, which can be carried by considering all values of I_{t+1} such that the pair $(\, I_t,\, I_{t+1})$ satisfies Inequality (5).

`// FORWARD PASS: Constructs the optimal sequence of` I_t`'s`

foreach $t \in (\,1,\,\ldots,\,T\,)$ **do**

  let $I_t^* = \omega_{\sigma,\lambda}^{[t-1]}(\, I_{t-1}\,)$

return $I^* \triangleq (\, I_1^*,\, \ldots,\, I_T^*\,)$

4. M-step: Estimation of the Sale and Loss Distribution Parameters

In this section we seek to compute values of the sale and loss distribution parameters $\sigma \in \mathbb{R}_{>0}$ and $\lambda \in [0,1]$, given the initial inventory level $I_0 \in \mathbb{Z}_{\geq 0}$, a feasible inventory level history $I \in \mathbb{Z}_{\geq 0}^T$, the sales history $S \in \mathbb{Z}_{\geq 0}^T$, and the replenishment history $R \in \mathbb{Z}_{\geq 0}^T$. To start, we note that in light of the arguments put forth in Section 3, this task is equivalent to that of maximizing function $\Phi_{\sigma,\lambda}$ over all feasible values of the parameters σ and λ. In turn, from Equations (8)-(10) we recognize that the aforementioned function is the sum of functions which depend only on the parameter σ with the sum of functions which depend only on the parameter λ. Indeed:

$$\Phi_{\sigma,\lambda}(\, I \,) \;=\; \underbrace{\left(\sum_{t=1}^{T} \log \mathbb{P}_{\sigma}(\, S_t \mid I_{t-1} \,) \right)}_{\triangleq\, \Phi_s(\sigma)} \;+\; \underbrace{\left(\sum_{t=1}^{T} \log \mathbb{P}_{\lambda}(\, I_t \mid I_{t-1},\, S_t,\, R_t \,) \right)}_{\triangleq\, \Phi_\ell(\lambda)} \tag{12}$$

Therefore, computing the optimal values of the parameters σ and λ can be readily carried by separately computing the maximizers of the single-argument functions Φ_s and Φ_ℓ over their respective domains, which in turn can be executed using any method for optimizing continuously differentiable nonlinear functions, *e.g.*, gradient ascent, Newton's Method, BFGS, etc.

5. Inventory Level Estimation

Once the sale and loss distribution parameters have been estimated, the probability distribution over inventory levels conditional on the initial inventory and on the sales and replenishments up to period t, *i.e.*,

$$\rho_t(\, I_t \,) \;\triangleq\; \mathbb{P}_{\sigma,\lambda}(\, I_t \mid I_0,\, S_1,\, \ldots,\, S_t,\, R_1,\, \ldots,\, R_t \,), \tag{13}$$

can be easily and efficiently computed. For this purpose, we first note that $\rho_1(\, I_1 \,) = \mathbb{P}_{\sigma,\lambda}(\, I_1 \mid I_0,\, S_1,\, R_1 \,)$ can be computed directly given the assumptions put forth in Section 2. Next, as noted in Ref. [5], the probability distributions associated with periods $t \geq 2$ can be sequentially computed according to the following recursion:

$$\forall\, I_t \in \{\, 0,\, \ldots,\, I_t^{max} \,\} :\; \rho_t(\, I_t \,) \;=\; \sum_{I_{t-1}=0}^{I_{t-1}^{max}} \rho_{t-1}(\, I_{t-1} \,)\, \mathbb{P}_{\sigma,\lambda}(\, I_t \mid I_{t-1},\, S_t,\, R_t \,) \tag{14}$$

Lastly, once we have computed $\rho_t(\, I_t \,)$, we can calculate the Marginal Maximum Likelihood Estimate (MMLE) of the inventory level conditional on all the information observed up to time t as:

$$I_t^{MMLE} \;\triangleq\; \underset{I_t \,\in\, \{\, 0,\, \ldots,\, I_t^{max} \,\}}{\arg\max} \; \rho_t(\, I_t \,) \tag{15}$$

6. Experiments and Results

In this section we experimentally evaluate the performance of our EM Algorithm by means of simple simulations. In particular, we simulate a 'naive' inventory management system which computes the current inventory as the previous inventory minus current sales plus current replenishments (*i.e.*, $I_t = I_{t-1} - S_t + R_t$ for each period t) and follows a $(\mathcal{Q},\, \mathcal{R}\,)$ policy. The system begins with an initial inventory level of $I_0 = 15$ units and it re-orders an amount of $\mathcal{Q} = 20$ units every time the inventory level reaches or falls below $\mathcal{R} = 10$ units. We simulate the naive system's evolution by feeding it random sales with the distributions described in Section 2 and parameter $\sigma = 5.0$, which is sampled considering the true physical inventory. Furthermore, we setup random losses with parameter $\lambda = 0.25$, which over a time horizon of $T = 60$ periods usually causes the naive system to freeze.

With this setup in place, we estimate the true physical inventory by running our EM Algorithm until the absolute changes of the estimates of σ and λ are within 0.01 units, respectively. Regarding the initial guesses of the parameters, we choose $\sigma^{[0]}$ to be the average of the sales during the experiment's horizon, as it seems like a reasonable choice, and we choose $\lambda^{[0]} = 0.5$, as it is the midpoint of the interval $[0, 1]$. The E-step is carried exactly as described in Section 3, while the M-step is carried approximately using 20 iterations of gradient ascent. Figure 2 shows the sale and inventory level histories for a typical simulation, where the inventory level according to the naive system is shown in red, the MMLE of the inventory level conditional on all previous observations (see Equation (15)) is shown in blue, and the true physical inventory level is shown in green.

Noteworthily, for most of the examples we simulated our EM Algorithm terminated after less than five iterations.

7. Future Research

There are several lines of future research which can stem from the work presented in this paper. For instance, the performance of our EM Algorithm should be evaluated on real data and compared to the performance of similar methods proposed in the literature. This would be interesting because in the real world the true physical inventory levels are not known. Another avenue of research is the extension of our model to one that can optimally decide on product re-orders and inventory inspections. This would be challenging because it leads to a Partially-observable

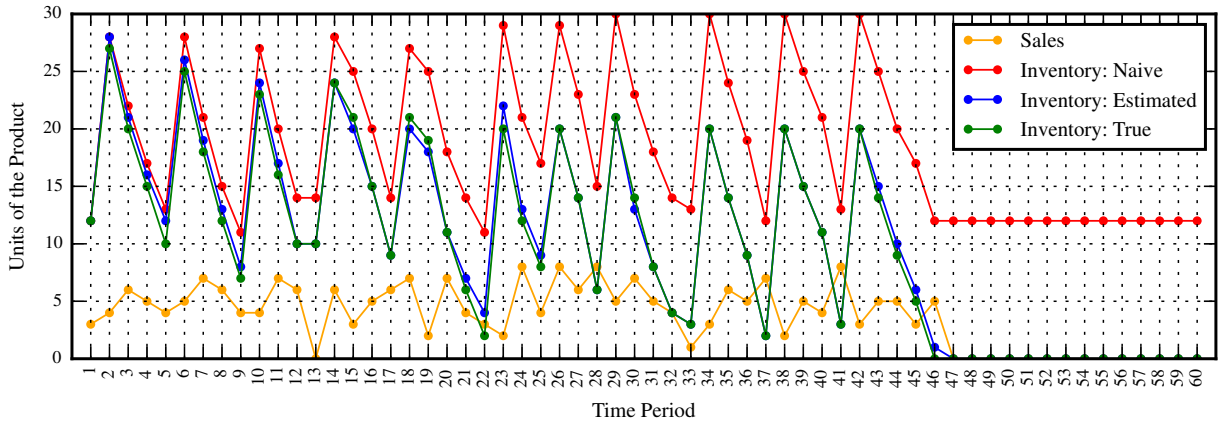

Figure 2: Sales and inventory time histories for a typical simulation. Here the initial guesses where $\sigma^{[0]} = 3.72$ and $\lambda^{[0]} = 0.50$, respectively. For this example, our EM Algorithm terminated after only two iterations, with estimates $\sigma^* = 5.20$ and $\lambda^* = 0.32$.

Markov Decision Process (POMDP) formulation, and its is well-known that even approximating the optimal policies for POMDPs can be, in general, computationally intractable.

Acknowledgements

The authors would like to thank the Board of Directors of Tiendas Industriales Asociadas Sociedad Anónima (TIA S.A.), a leading grocery retailer in Ecuador, for authorizing their company to provide the authors with time history data of sales, replenishments and inventory levels for hundreds of their products and several of their stores. Explorative analysis of this data lead to the discovery of products which seem to have suffered from stock freezing, which in turn lead to the design of our DBN model and our EM algorithm.

References

[1] Yun Kang and Stanley B Gershwin. Information Inaccuracy in Inventory Systems: Stock Loss and Stockout. *IIE Transactions*, 37(9):843–859, 2005.

[2] Donald L Iglehart and Richard C Morey. Inventory Systems with Imperfect Asset Information. *Management Science*, 18(8):B–388, 1972.

[3] Ananth Raman, Nicole DeHoratius, and Zeynep Ton. Execution: The Missing Link in Retail Operations. *California Management Review*, 43(3):136–152, 2001.

[4] Nicole DeHoratius and Ananth Raman. Inventory Record Inaccuracy: An Empirical Analysis. *Management Science*, 54(4):627–641, 2008.

[5] Nicole DeHoratius, Adam J Mersereau, and Linus Schrage. Retail Inventory Management when Records are Inaccurate. *Manufacturing & Service Operations Management*, 10(2):257–277, 2008.

[6] Li Chen and Adam J Mersereau. Analytics for Operational Visibility in the Retail Store: The Cases of Censored Demand and Inventory Record Inaccuracy. In *Retail Supply Chain Management*, pages 79–112. Springer, 2015.

[7] Christopher M Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[8] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.